PLMIX: An R package for modeling and clustering partially ranked data

Cristina Mollica
Department of Statistics
Sapienza University of Rome

Luca Tardella
Department of Statistics
Sapienza University of Rome

arXiv:1612.08141v2 [stat.CO] 9 Jan 2017

Abstract

Ranking data represent a peculiar form of multivariate ordinal data taking values in the set of permutations. Despite the numerous methodological contributions to increase the flexibility of ranked data modeling, the application of more sophisticated models is limited by the related computational issues. The **PLMIX** package offers a comprehensive framework aiming at endowing the R environment with the recent methodological advancements. The usefulness of the novel **PLMIX** package can be motivated from several perspectives: (i) it contributes to fill the gap concerning the Bayesian estimation of ranking models in R, by focusing on the *Plackett-Luce* model as generative distribution and its extension within the finite mixture approach; (ii) it combines the flexibility of R routines and the speed of compiled C code, with possible parallel execution; (iii) it covers the fundamental phases of ranking data analysis, allowing for a more careful and critical application of ranking models in real experiments: (iv) it provides effective tools for clustering heterogeneous partially ranked data. The functionality of the novel package is illustrated with application to simulated and real datasets.

Keywords: Ranking data, Plackett-Luce model, Mixture models, Bayesian inference, MAP estimation, MCMC methods, R.

Contents

1. Introduction

A *ranking* is an ordered sequence resulting from the comparative evaluation of a given set of *items*. This type of data is common in a large number of research fields, as testified by

the widespread applications of ranking data analysis. For example, various political election systems allow voters to express multiple preferences (not only the most-liked candidate) in the ballots (Diaconis 1987; Stern 1993; Gormley and Murphy 2008a,b) and a method to aggregate the resulting rankings into a final *consensus* is needed, as in Davenport and Kalagnanam (2004); Meilă, Phadnis, and Patterson (2007); Bargagliotti (2009). Furthermore, marketing and psychological studies are typically aimed at investigating individual preferences (Vigneau, Courcoux, and Semenou 1999; Gormley and Murphy 2010; Vitelli, Sørensen, Frigessi, and Arjas 2014) and attitudes (Yu, Lam, and Lo 2005; Gormley and Murphy 2006) towards multiple alternatives.

Mollica and Tardella (2014) proposed the conversion of the quantitative multivariate profiles resulting from a bioassay experiment into ranking sequences. The ranking transformation was motivated as a possible data normalization method in the case when a well-established pre-processing technique is lacking. In their work, the ranking-based analysis was proved to provide a more robust evidence to successfully characterize the disease status.

Of course, also sports and competitions naturally contribute to motivate the development of ranking analysis, in order to quantify the ability of the competitors from the ordinal evidence, see for example Henery (1981); Stern (1990); Clerides and Stengos (2012); Caron and Doucet (2012). Further references can be found in the the recent monograph on ranking data analysis offered by Alvo and Yu (2014). A public data repository devoted to preference data, although not necessarily rankings, is available at `http://www.preflib.org/` (Mattei and Walsh 2013).

From a mathematical perspective, ranking data represent a well-characterized form of multivariate ordinal data, with a direct correspondence with the set of permutations (Plackett 1968; Henery 1983; Stern 1990). By following the reviews in Critchlow, Fligner, and Verducci (1991); Marden (1995); Alvo and Yu (2014), the main approaches to conceive parametric ranking models can be classified in four classes:

- *order statistics models* (OS), also known as random utility models, whose cornerstone is the *Thurstone model* (Thurstone 1927);
- *paired comparison models*, where the most general family is given by the Babington-Smith model described in Babington Smith (1950);
- *distance-based models* (DB), originally introduced by Mallows (1957) and also referred to as Mallows' models;
- *stagewise models*.

Each model class alludes to a specific generative process of the ordinal judgment on the given set of alternatives. This work concentrates on the last parametric family bearing on the decomposition of the ranking process into consecutive stages, that is, the sequential selections of the items in order of preference. In particular, our interest is in the *Plackett-Luce model* (PL) and its extensions by assuming the Bayesian inferential perspective.

Despite the numerous methodological contributions of the last decades, enhancing the flexibility of the aforementioned model classes, the application of more sophisticated ranking models is still limited in practice. Likely, the main reason lies in the computational complexity emerging from the peculiar multivariate structure of ranking data, requiring the development of specialized software. This might have slowed down a wider use of the most recent model

Table 1: Characteristics of the existing R packages for ranking data compared with the novel **PLMIX** package.

Package	Ranking type	Model class	Mixture	Inference
PerMallows	Complete	DB, GMM	No	MLE
pmr	Complete	DB, WDB, PL	No	MLE
prefmod	Partial	BT	Yes	MLE
Rankcluster	Partial	ISR	Yes	MLE
rankdist	Partial	DB	Yes	MLE
RMallow	Partial	DB	Yes	MLE
StatMethRank	Complete	MNOS, WDB	Yes	MLE*
StatRank	Partial	OS	Yes	MLE
PLMIX	Partial	PL	Yes	Bayesian**

* a single function to fit the Bayesian MNOS is available
** MLE can be recovered as special case under noninformative prior setting

proposals. The **PLMIX** package offers a comprehensive framework aiming at enriching the R environment with one of the recent methodological advancements in model-based ranked data analysis, by adequately accounting for the related computational issues.

The paper is organized as follows: Section 2 provides a detailed review of the existing packages in R for the analysis of ranking data and highlights the main differences with the novel **PLMIX** package. The PL is briefly reviewed in Section 3 and its Bayesian extension to the finite mixture setting is detailed in Section 4 with the related inferential procedures. Section 5 describes the implementation of the functions included in the **PLMIX** package with some simulated and real data examples. The paper ends in Section 6 with some remarks and suggestions for future developments.

2. Review of R packages for ranking data analysis

Several R packages are currently available for model-based analysis of ranking data and their main features are summarized in Table 1, in comparison with the novel **PLMIX** package. Some essential features of each package are provided in the following:

- **PerMallows**, described in Irurozki, Calvo Molinos, and Lozano Alonso (2014), offers a suite of functions for the MLE of DBs and their multiparametric extensions, referred to as *Generalized Mallows models* (GMM) in the seminal work Fligner and Verducci (1986). Various metrics for the ranking space are considered but partial rankings and finite mixtures are not contemplated;

- **pmr**, presented in Lee and Yu (2013), applies standard MLE methods to infer several parametric ranking models, such as the DB, the *weighted distance-based models* (WDB) proposed by (Lee and Yu 2010) and the PL. Ranking models are considered in their basic form (no mixture) and only complete rankings are allowed;

- **prefmod**, introduced in Hatzinger and Dittrich (2012), considers the equivalent repre-

sentation of the rankings in the form of pairwise comparisons and the application of the *Bradley-Terry* model (BT), see Bradley and Terry (1952); Bradley (1984). This package allows for the inclusion of individual and item-specific covariates, the handling of partial rankings and the generalization to the finite mixture setting, addressed in the MLE and nonparametric MLE framework;

- **rankdist**, relying on the EM algorithm to fit mixtures of DBs with various metrics; it accepts both complete and partial rankings, see Murphy and Martin (2003) for the methodological contribution;

- **RMallow** implements the mixture of DBs with the Kendall distance as metric on the ranking space. Both complete and partial rankings are allowed;

- **Rankcluster** is widely described in Jacques, Grimonprez, and Biernacki (2014). It implements the mixture of *Insertion Sort Rank data models* (ISR) proposed by Jacques and Biernacki (2014), which extends the generative model for ranking data based on the Insertion Sort algorithm, originally defined in Biernacki and Jacques (2013), into the mixture approach. The ISR mixture is motivated as model-based clustering tool of partial and potentially multivariate (hierarchical) ranking data;

- **StatMethRank** supports the recent review of the ranking literature provided by Alvo and Yu (2014). Regarding model-based analysis, it implements the Bayesian Multivariate Normal ordered statistics model (MNOS) described in Yu (2000) and the mixture of WDBs from the MLE framework, but exclusively on complete rankings.

- **StatRank** covers the class of random utility models involving the PL as special instance. Estimation is carried on by means of the Generalized Method-of-Moments (Soufiani 2014; Soufiani, Chen, Parkes, and Xia 2013a) and can be performed also on partial observations. A function to fit the mixture of random utility models is also provided (Soufiani, Diao, Lai, and Parkes 2013b).

The comparison with the available packages in R (Table 1) points out that the existing contributed libraries cover a wide range of the parametric options reviewed in Section 1. Moreover, most of them account also for the possible presence of incomplete observations and for the generalization of the ranking generative mechanism to the mixture framework.

Nevertheless, with the only exception of the function `mvnos.model` in the **StatMethRank** package which estimates the Bayesian MNOS model via MCMC methods, only the novel **PLMIX** is specifically devoted to the Bayesian inference for ranking data. In this setting, the MLE approach is recovered as special case of the Bayesian analysis with a noninformative prior specification. More specifically, **PLMIX** performs Bayesian estimation of ranking models in R, by focusing on the Plackett-Luce model and its finite mixture extensions as sampling distribution. To address the issue of computationally demanding procedures, **PLMIX** can take advantage of a hybrid code linking the R environment with the C++ programming language. The parallelization option is also implemented, such that finite mixtures with a different number of components can be simultaneously analyzed by the user. In addition to inferential techniques, the package assists other fundamental phases related to the mixture model analysis, such that the selection of the optimal number of components and the goodness-of-fit assessment, aimed at a more critical clustering of partial rankings.

3. The Plackett-Luce model for partial orderings

3.1. Preliminaries and data format

Let us first clarify the basic terminology for the data input, in particular the difference between ranking and ordering. Formally, a *full* (or *complete*) *ranking* $\pi : I \to R$ is a bijective mapping of a finite set $I = \{1, \ldots, K\}$ of labeled *items* (or alternatives) into a set of *ranks* $R = \{1, \ldots, K\}$, resulting from the attribution of a rank to each item according to a determined criterion. Hence, K is the total number of items to be ranked and the result of the mapping can be represented in terms of the K-tuple $\pi = (\pi(1), \ldots, \pi(K))$, where the generic entry $\pi(i)$ indicates the rank assigned to the i-th item. According to the underlying convention in the ranking theory, if $\pi(i) < \pi(i')$ then item i is said to be ranked higher than/preferred to item i'.

Ranking data admit an alternative representation in terms of orderings. Specifically, the *full* (or *complete*) *ordering* $\pi^{-1} : R \to I$ is simply the inverse function of the ranking π, yielding the ordered vector $\pi^{-1} = (\pi^{-1}(1), \ldots, \pi^{-1}(K))$ whose generic component $\pi^{-1}(j)$ denotes the item ranked in the j-th position.

In many real applications, for example when K is large, the ranking elicitation could be not completely carried out. A typical situation is when the ranker specifies only her most-liked $t < K$ items and leaves the remaining $K - t$ positions undefined. In this case the generic observation consists in the so-called *top-t partial ordering* of the form $\pi^{-1} = (\pi^{-1}(1), \ldots, \pi^{-1}(t))$. With a slight abuse of notation, the remaining $K - t$ alternatives are tacitly assumed to be ranked lower, formally $\pi(i) > t$ for all $i \notin \{\pi^{-1}(1), \ldots, \pi^{-1}(t)\}$. Notice that a complete ordering is a special instance of top-t partial ordering with $t = K - 1$. In fact, a top-$(K-1)$ sequence corresponds to a complete ordering, since the single missing K-th entry can be unambiguously determined.

Finally, we remark that in the present context ties, i.e., the case when multiple items occupy the same position, are not contemplated.

3.2. The Plackett-Luce model

The Plackett-Luce model, whose paternity is jointly attributed to Luce (1959) and Plackett (1975), is one of the most successfully applied parametric distribution to model partially ranked data. It belongs to the class of *stagewise ranking models*, where the elicitation of the comparative judgment on the K alternatives is conceived as a sequential process divided into $K - 1$ stages: similarly to a random draw without replacement from an urn, at each stage the most-liked item is specified among the alternatives not selected at the previous stages. The sequential random selections of the items are governed by the *support parameters* $\underline{p} = (p_1, \ldots, p_K)$, that is, positive constants representing a measure of liking toward each item. More specifically, let $\pi_s^{-1} = (\pi_s^{-1}(1), \ldots, \pi_s^{-1}(n_s))$ be a generic partial ordering, where n_s is the number of items ranked by unit s in the first n_s positions. The PL postulates

$$\mathbf{P}_{\mathrm{PL}}(\pi_s^{-1}|\underline{p}) = \prod_{t=1}^{n_s} \frac{p_{\pi_s^{-1}(t)}}{\sum_{i=1}^{K} p_i - \sum_{\nu=1}^{t-1} p_{\pi_s^{-1}(\nu)}}. \tag{1}$$

As apparent in (1), the underlying assumption is that the relative preference of each item over the other ones is constant across the ranking stages. For a given a random sample

$\underline{\pi}^{-1} = \{\pi_s^{-1}\}_{s=1}^N$ of N partial top orderings, the observed-data log-likelihood is equal to

$$l(\underline{p}) = \sum_{i=1}^{K} \gamma_i \log p_i - \sum_{s=1}^{N} \sum_{t=1}^{n_s} \log \sum_{i=1}^{K} \delta_{sti} p_i, \tag{2}$$

where $\gamma_i = \sum_{s=1}^N u_{si}$ with $u_{si} = \mathbb{I}_{[i \in \{\pi_s^{-1}(1),\ldots,\pi_s^{-1}(n_s)\}]}$ and $\delta_{sti} = \mathbb{I}_{[i \notin \{\pi_s^{-1}(1),\ldots,\pi_s^{-1}(t-1)\}]}$ with $\delta_{s1i} = 1$ for all $s = 1, \ldots, N$ and $i = 1, \ldots, K$.

4. Bayesian inference of the Plackett-Luce mixture model

In this section we give a brief outline of the Bayesian approach based on the data augmentation strategy to make inference on the PL parameters, both in the case of homogeneous population without an underlying group structure and in the more general finite mixture framework.

4.1. The homogeneous case

Because of the normalization term $\sum_{i=1}^K \delta_{sti} p_i$ in (2), the direct maximization of the log-likelihood is not straightforward. In the Bayesian setting, simple and effective estimation procedures for the PL were introduced by Caron and Doucet (2012). Their crucial idea to overcome the inconvenient form of the PL likelihood is a data augmentation step with continuous latent variables associated to each entry of the observed matrix. More specifically, Caron and Doucet (2012) suggest to employ auxiliary variables $\underline{y} = (y_{st})$ for $s = 1, \ldots, N$ and $t = 1, \ldots, n_s$ with a suitable parametric assumption for their joint conditional distribution, given by

$$f(\underline{y} | \underline{\pi}^{-1}, \underline{p}) = \prod_{s=1}^{N} \prod_{t=1}^{n_s} f_{\text{Exp}}\left(y_{st} \middle| \sum_{i=1}^{K} \delta_{sti} p_i \right), \tag{3}$$

where $f_{\text{Exp}}(\cdot|\lambda)$ is the Negative Exponential density function indexed by the rate parameter λ. Assumption (3) implies the following complete-data likelihood

$$L_{\text{c}}(\underline{p}, \underline{y}) = \prod_{i=1}^{K} p_i^{\gamma_i} e^{-p_i \sum_{s=1}^N \sum_{t=1}^{n_s} \delta_{sti} y_{st}}, \tag{4}$$

which clearly suggests the adoption of independent Gamma priors for the support parameters $\underline{p}$ to obtain a conjugate model specification. Formally, the conjugacy of the Gamma density with model (1) is explained by the Thurstonian interpretation of the PL as OS with the latent item scores distributed according to independent Gumbel distributions, see Yellott (1977) for further details.

MAP estimation via EM algorithm

As first type of inference, Caron and Doucet (2012) describe how to achieve the Maximum A Posteriori (MAP) estimate of the vector $\underline{p}$, i.e., the mode of the posterior distribution. In the presence of latent variables, the popular EM algorithm introduced by Dempster, Laird, and Rubin (1977) can be applied to optimize the posterior distribution where, differently from the canonical frequentist implementation, the objective function contemplates also the extra information provided by the prior distribution $f_0(\underline{p}) = \prod_{i=1}^K f_{\text{Ga}}(p_i)$. Explicitly, by denoting

with $l_c(\underline{p}, \underline{y}) = \log L_c(\underline{p}, \underline{y})$ the complete-data log-likelihood, the E-step returns this expression of the objective function

$$\begin{aligned} Q(\underline{p}, \underline{p}^*) &= \mathbb{E}[l_c(\underline{p}, \underline{y}) | \, \underline{\pi}^{-1}, \underline{p}^*] + \log f_0(\underline{p}) \\ &\propto \sum_{i=1}^{K} \gamma_i \log p_i - \sum_{s=1}^{N} \sum_{t=1}^{n_s} \frac{\sum_{i=1}^{K} \delta_{sti} p_i}{\sum_{i=1}^{K} \delta_{sti} p_i^*} + \sum_{i=1}^{K} ((c-1) \log p_i - d p_i), \end{aligned} \tag{5}$$

where c and d are, respectively, the shape and rate parameter of the Gamma prior. The direct maximization of (5) yields a closed-form solution for the iterative update of the PL parameters that converges to the MAP estimate of $\underline{p}$. At the generic iteration $l+1$, the update of the M-step proceeds as follows

$$p_i^{(l+1)} = \frac{c - 1 + \gamma_i}{d + \sum_{s=1}^{N} \sum_{t=1}^{n_s} \dfrac{\delta_{sti}}{\sum_{i=1}^{K} \delta_{sti} p_i^{(l)}}} \qquad i = 1, \ldots, K.$$

Notice that, by setting noninformative hyperparameters $c = 1$ and $d = 0$, the M-step reduces to the iterative Minorization Maximization algorithm described by Hunter (2004) to compute the MLE of the PL.

Gibbs sampling

Caron and Doucet (2012) describe also the GS procedure, a simulation-based method to approximate the joint posterior distribution and to assess the uncertainty of the parameter estimates with empirical summaries of posterior variability. In this regard, first notice that the full-conditional of $\underline{y}$ is imposed by the data augmentation assumption (3). Finally, thanks to the conjugate prior specification, the full-conditionals of the support parameters still belong to the Gamma family. Specifically, at the generic iteration $l+1$ the GS alternates the following two simulation steps

$$\begin{aligned} y_{st}^{(l+1)} | \pi_s^{-1}, \underline{p}^{(l)} &\sim \text{Exp}\left(\sum_{i=1}^{K} \delta_{sti} p_i^{(l)} \right), \\ p_i^{(l+1)} | \underline{\pi}^{-1}, \underline{y}^{(l+1)} &\sim \text{Ga}\left(c + \gamma_i, d + \sum_{s=1}^{N} \sum_{t=1}^{n_s} \delta_{sti} y_{st}^{(l+1)} \right). \end{aligned}$$

4.2. The finite PL mixture

We now review the proposal recently developed by Mollica and Tardella (2016, published online) to extend the data augmentation approach (3) to the finite mixture context.

Formally, the G-component PL mixture model assumes that observations are sampled from a heterogeneous population composed of G subpopulations, called *mixture components*. Each component g follows a basic PL distribution with a specific support parameter vector $\underline{p}_g$, that is,

$$\pi_s^{-1} | \underline{p}, \underline{\omega} \sim \sum_{g=1}^{G} \omega_g \, \mathbf{P}_{\text{PL}}(\pi_s^{-1} | \underline{p}_g), \tag{6}$$

where $\underline{\omega} = (\omega_1, \ldots, \omega_G)$ are referred to as *mixture weights*. Let $\underline{z}_s = (z_{s1}, \ldots, z_{sG}) | \underline{\omega} \sim \text{Multinom}(1, \underline{\omega} = (\omega_1, \ldots, \omega_G))$ be the vector describing the latent group membership of unit s, whose generic entry is

$$z_{sg} = \begin{cases} 1 & \text{if unit } s \text{ belongs to the } g\text{-th mixture component,} \\ 0 & \text{otherwise.} \end{cases}$$

As apparent, the parametric representation (6) defines a latent variable model where the unobserved group configuration is marginalized out. Mollica and Tardella (2016, published online) suggest to account for the latent group structure in the data augmentation step as follows

$$\begin{aligned} \underline{\omega} &\sim \text{Dir}(\alpha_1, \ldots, \alpha_G) \\ p_{gi} &\overset{i}{\sim} \text{Ga}(c_{gi}, d_g) \\ \underline{z}_s | \, \underline{\omega} &\overset{iid}{\sim} \text{Multinom}(1, \underline{\omega}) \\ \pi_s^{-1} | \underline{z}_s, \underline{p} &\overset{i}{\sim} \prod_{g=1}^{G} \mathbf{P}_{\text{PL}}(\pi_s^{-1} | \underline{p}_g)^{z_{sg}} \\ y_{st} | \pi_s^{-1}, \underline{z}_s, \underline{p} &\overset{i}{\sim} \text{Exp}\left(\prod_{g=1}^{G} \left(\sum_{i=1}^{K} \delta_{sti} p_{gi} \right)^{z_{sg}} \right). \end{aligned}$$

The resulting complete-data likelihood is given by

$$\begin{aligned} L_{\text{c}}(\underline{p}, \underline{\omega}, \underline{y}, \underline{z}) &= \mathbf{P}(\underline{y}, \underline{\pi}^{-1}, \underline{z} | \underline{p}, \underline{\omega}) = \prod_{s=1}^{N} \mathbf{P}(\underline{y}_s, \underline{\pi}_s^{-1}, \underline{z}_s | \underline{p}, \underline{\omega}) \, \mathbf{P}(\underline{\pi}_s^{-1} \, | , \underline{z}_s, \underline{p}) \, \mathbf{P}(\underline{z}_s | \underline{p}) \\ &= \prod_{s=1}^{N} \prod_{g=1}^{G} \left(\omega_g \prod_{i=1}^{K} p_{gi}^{u_{si}} e^{-p_{gi} \sum_{t=1}^{n_s} \delta_{sti} y_{st}} \right)^{z_{sg}}. \end{aligned}$$

Similarly to Caron and Doucet (2012), also Mollica and Tardella (2016, published online) detail the MAP estimation and the GS algorithm to perform a fully Bayesian inference on the finite PL mixture.

MAP estimation via EM algorithm

In the mixture setting, the EM algorithm alternates the E-step, which requires the computation of $Q((\underline{p}, \underline{\omega}), (\underline{p}^*, \underline{\omega}^*)) = \mathbb{E}_{\underline{y}, \underline{z} | \, \underline{\pi}^{-1}, \underline{p}^*, \underline{\omega}^*}[l_c(\underline{p}, \underline{\omega}, \underline{y}, \underline{z})] + \log f_0(\underline{p}, \underline{\omega})$ w.r.t. the conditional distribution of all the latent variables, and the subsequent maximization, yielding the follow-

ing three updates for the M-step at the generic iteration $(l+1)$:

$$\hat{z}_{sg}^{(l+1)} = \frac{\omega_g^{(l)} \, \mathbf{P}_{\text{PL}}(\pi_s^{-1}|\underline{p}_g^{(l)})}{\sum_{g'=1}^{G} \omega_{g'}^{(l)} \, \mathbf{P}_{\text{PL}}(\pi_s^{-1}|\underline{p}_{g'}^{(l)})},$$

$$\omega_g^{(l+1)} = \frac{\alpha_g - 1 + \sum_{s=1}^{N} \hat{z}_{sg}^{(l+1)}}{\sum_{g'=1}^{G} \alpha_{g'} - G + N},$$

$$p_{gi}^{(l+1)} = \frac{c_{gi} - 1 + \hat{\gamma}_{gi}^{(l+1)}}{d_g + \sum_{s=1}^{N} \hat{z}_{sg}^{(l+1)} \sum_{t=1}^{n_s} \frac{\delta_{sti}}{\sum_{i=1}^{K} \delta_{sti} p_{gi}^{(l)}}},$$

where $\hat{\gamma}_{gi}^{(l+1)} = \sum_{s=1}^{N} \hat{z}_{sg}^{(l+1)} u_{si}$. Interestingly, with the noninformative prior setting, corresponding to $c_{gi} = 1$, $d_g = 0$ and $\alpha_g = 1$, in the above MAP procedure one can recover the MLE of the PL mixture presented by Gormley and Murphy (2006).

Obviously, the MAP solution offers a suitable starting point to initialize the GS algorithm as will be illustrated in the next subsection.

Gibbs sampling

Thanks to the conjugate prior specification, all the full-conditional distributions have known form and are easy to be sampled. At the generic iteration $l+1$, the GS algorithm consists in iteratively generating random values from the following full-conditionals

$$\begin{aligned}
\underline{\omega}^{(l+1)}|\underline{z}^{(l)} &\sim \text{Dir}\left(\alpha_1 + \sum_{s=1}^{N} z_{s1}^{(l)}, \ldots, \alpha_G + \sum_{s=1}^{N} z_{sG}^{(l)}\right) \\
y_{st}^{(l+1)}|\pi_s^{-1}, \underline{z}_s^{(l)}, \underline{p}^{(l)} &\sim \text{Exp}\left(\prod_{g=1}^{G}\left(\sum_{i=1}^{K} \delta_{sti} p_{gi}^{(l)}\right)^{z_{sg}^{(l)}}\right) \\
p_{gi}^{(l+1)}|\underline{\pi}^{-1}, \underline{y}^{(l+1)}, \underline{z}^{(l)} &\sim \text{Gam}\left(c_{gi} + \gamma_{gi}^{(l)}, d_g + \sum_{s=1}^{N} z_{sg}^{(l)} \sum_{t=1}^{n_s} \delta_{sti} y_{st}^{(l+1)}\right) \\
\underline{z}_s^{(l+1)}|\pi_s^{-1}, \underline{y}_s^{(l+1)}, \underline{p}^{(l+1)}, \underline{\omega}^{(l+1)} &\sim \text{Multinom}\left(1, \left(m_{s1}^{(l+1)}, \ldots, m_{sG}^{(l+1)}\right)\right),
\end{aligned}$$

where $\gamma_{gi}^{(l)} = \sum_{s=1}^{N} z_{sg}^{(l)} u_{si}$ and

$$m_{sg}^{(l+1)} = \frac{\omega_g^{(l+1)} \prod_{i=1}^{K} (p_{gi}^{(l+1)})^{u_{si}} e^{-p_{gi}^{(l+1)} \sum_{t=1}^{n_s} \delta_{sti} y_{st}^{(l+1)}}}{\sum_{g'=1}^{G} \omega_{g'}^{(l+1)} \prod_{i=1}^{K} (p_{g'i}^{(l+1)})^{u_{si}} e^{-p_{g'i}^{(l+1)} \sum_{t=1}^{n_s} \delta_{sti} y_{st}^{(l+1)}}},$$

see Mollica and Tardella (2016, published online) for more analytical details.

Table 2: Model selection criteria implemented in the **PLMIX** package.

Criterion	Formula
DIC_1	$\bar{D} + (\bar{D} - D(\hat{\theta}_{\text{MAP}}))$
DIC_2	$\bar{D} + \frac{\mathbb{VAR}[D(\theta)\vert\underline{\pi}^{-1}]}{2}$
BPIC_1	$\bar{D} + 2(\bar{D} - D(\hat{\theta}_{\text{MAP}}))$
BPIC_2	$\bar{D} + 2\mathbb{VAR}[D(\theta)\vert\underline{\pi}^{-1}]$
BICM_1	$\bar{D} + \frac{\mathbb{VAR}[D(\theta)\vert\underline{\pi}^{-1}]}{2}(\log N - 1)$
BICM_2	$D(\hat{\theta}_{\text{MAP}}) + \frac{\mathbb{VAR}[D(\theta)\vert\underline{\pi}^{-1}]}{2}\log N$

Label-switching issue

The *label switching* (LS) is an identifiability issue that can hamper the straightforward use of the MCMC simulations for the Bayesian estimation of mixture models (Marin, Mengersen, and Robert 2005). It reflects the arbitrary attribution of the indices $\{1, \ldots, G\}$ to denote the mixture components, such that the relabeling of the latent classes does not modify the resulting sampling distribution. If the likelihood invariance is fulfilled also by the prior distribution, the marginal posterior distribution of each parameter is the same for all the mixture components, nullifying the standard practice to use the posterior summaries as point estimates.

Among the several strategies proposed in the statistical literature to solve the LS issue in MCMC analysis of mixture models, see Jasra, Holmes, and Stephens (2005) for a review, we focus on the relabeling algorithms (RA). The basic idea is the *ex-post* relabeling of the raw MCMC samples in order to derive meaningful posterior estimates. A comprehensive review of the RA can be found in Papastamoulis (2016), describing their implementation in the recently released R package **label.switching**. We have exploited this package for the treatment of the LS in our Bayesian PL mixture applications.

4.3. Bayesian model comparison criteria

A crucial step in the finite mixture analysis is the determination of the optimal number $\hat{G}$ of components which, in general, is not known *a priori*. For a critical up-to-date review of Bayesian model selection tools, see Gelman, Hwang, and Vehtari (2014).

The **PLMIX** package includes several Bayesian model selection criteria to compare PL mixture models with a different number of components fitted on the same data set. The considered measures are:

- *Deviance Information Criterion* (DIC);
- *Bayesian Predictive Information Criterion* (BPIC);
- *Bayesian Information Criterion-Monte Carlo* (BICM);

whose analytic expressions are reported in Table 2. More specifically, let $D(\theta) = -2\log L(\theta)$ be the *deviance function* and $\bar{D} = \mathbb{E}[D(\theta)|\underline{\pi}^{-1}]$ its posterior expectation. In the original

formula introduced by Spiegelhalter, Best, Carlin, and Van Der Linde (2002),

$$\text{DIC} = \bar{D} + p_D$$

two alternative ways of computing of the *effective number of parameters* (p_D) are implemented: in DIC_1 the penalty term is given by the difference between $\bar{D}$ and the deviance evaluated in the MAP estimate $\hat{\theta}_{\text{MAP}}$, whereas in DIC_2 the model complexity is computed as half the posterior variance of the deviance $\mathbb{VAR}[D(\theta)|\underline{\pi}^{-1}]$, following the suggestion by Gelman, Carlin, Stern, and Rubin (2004). The latter DIC definition corresponds to the AICM, that is, the Bayesian version of AIC presented by Raftery, Satagopan, Newton, and Krivitsky (2007).

With the aim at correcting DIC's tendency to privilege more complex models (overfitting), Ando (2007) suggested to doubling the penalty term of DIC, resulting in the two measures BPIC_1 and BPIC_2 that parallel, respectively, DIC_1 and DIC_2.

BICM_1 and BICM_2, instead, represent the two Bayesian variants of BIC: the former is completely based on the posterior MCMC sample, as described in Raftery *et al.* (2007), whereas in the latter we employ the MAP estimate derived from the EM algorithm instead of its approximation from the GS output.

We advocate the use of $\hat{\theta}_{\text{MAP}}$ as point estimate for the mixture model parameters, instead of the posterior mean $\mathbb{E}[\theta|\underline{\pi}^{-1}]$ since the MAP estimates

- straightforwardly provides a meaningful estimate not affected by the LS problem;
- guarantees a positive value of model complexity p_D
- coincides with the MLE solution $\hat{\theta}_{\text{MLE}}$ under uninformative prior specification.

It follows that, given the likelihood invariance described in 4.2.3, all the considered model comparison measures do not suffer of the presence of LS. Thus, their estimation does not require the preliminary relabeling of the MCMC output.

4.4. Bayesian model assessment

Evaluating the fitting performance of a parametric model can be less straightforward in ranking data applications than in other multivariate contexts. In the frequentist domain, for example, model assessment is typically addressed with the computation of the *p-value* associated to a goodness-of-fit statistic, such as the likelihood ratio or Pearson's X^2. However, in sparse data situations characterized by a small sample size N w.r.t. the cardinality $K!$ of the ranking space, serious issues arise with this approach, since the chi-square distribution of the above statistics under the posited model H no longer applies. Cohen and Mallows (1983) suggested to overcome this difficulty by comparing observed and expected frequencies regarding relevant partitions of the ranking space.

The approach relying on the identification of meaningful groupings of the rankings has been successfully applied also within the Bayesian paradigm, where it is possible to generalize the classical test statistic into a parameter-dependent quantity, referred to as *discrepancy variable* (Gelman, Meng, and Stern 1996; Meng 1994). In order to assess the adequacy of the Bayesian PL mixture, the **PLMIX** package provides diagnostic tools derived from two significant summary statistics in ranked data analysis:

1. the most-liked item frequency vector $\underline{r}(\underline{\pi}^{-1})$, whose generic entry is

$$r_i(\underline{\pi}^{-1}) = \sum_{s=1}^{N} I_{[\pi_s^{-1}(1)=i]}$$

corresponding to the number of times that item i is ranked first;

2. the paired comparison frequency matrix $\tau(\underline{\pi}^{-1})$, whose generic entry is

$$\tau_{ii'}(\underline{\pi}^{-1}) = \sum_{s=1}^{N} (u_{si} + u_{si'} - u_{si}u_{si'}) I_{[\pi_s(i)<\pi_s(i')]}$$

corresponding to the number of times that item i is preferred to item i'.

As described in Mollica and Tardella (2016, published online), one could employ the aforementioned sample quantities to define the corresponding chi-square discrepancies as follows

$$X^2_{(1)}(\underline{\pi}^{-1};\theta) = \sum_{i=1}^{K} \frac{\left(r_i(\underline{\pi}^{-1}) - r_i^*(\theta)\right)^2}{r_i^*(\theta)} \qquad X^2_{(2)}(\underline{\pi}^{-1};\theta) = \sum_{i<i'} \frac{(\tau_{ii'}(\underline{\pi}^{-1}) - \tau^*_{ii'}(\theta))^2}{\tau^*_{ii'}(\theta)}. \tag{7}$$

The symbol * denotes the theoretical frequencies expected under the PL mixture model with parameter vector $\theta = (\underline{p}, \underline{\omega})$. The theoretical frequencies are, respectively,

$$r_i^*(\theta) = N p_i \qquad \text{and} \qquad \tau^*_{ii'}(\theta) = T_{ii'} \frac{p_i}{p_i + p_{i'}},$$

where $p_i = \sum_{g=1}^{G} \omega_g p_{gi}$ is the marginal overall support to item i and $T_{ii'} = \tau_{ii'} + \tau_{i'i}$ is the total number of pairwise comparisons between item i and i'. The reader can refer to Cohen and Mallows (1983), Yao and Böckenholt (1999), Yu (2000) and Gormley and Murphy (2009) for the expression of the above expected frequencies under different model specifications. For a given discrepancy variable $X^2(\underline{\pi}^{-1};\theta)$, the posterior predictive check of model goodness-of-fit relies on the computation of the *posterior predictive p-value*, corresponding to

$$p_B = \mathbf{P}(X^2(\underline{\pi}^{-1}_{\text{rep}};\theta) \geq X^2(\underline{\pi}^{-1}_{\text{obs}};\theta) | \, \underline{\pi}^{-1}_{\text{obs}}, H). \tag{8}$$

As emphasized by Gelman *et al.* (1996), once an MCMC sample of size L from the posterior distribution is available, the estimation of (8) can be obtained by iterating the following two-step procedure for all the parameter draws $\theta^{(l)}$: (i) simulate a new data set $\underline{\pi}^{-1(l)}_{\text{rep}}$ from the sampling distribution induced by $\theta^{(l)}$; (ii) compare $X^2(\underline{\pi}^{-1(l)}_{\text{rep}};\theta^{(l)})$ with $X^2(\underline{\pi}^{-1}_{\text{obs}};\theta^{(l)})$. The posterior predictive p-value (8) is finally estimated as the proportion of the times that $X^2(\underline{\pi}^{-1(l)}_{\text{rep}};\theta^{(l)})$ exceeds $X^2(\underline{\pi}^{-1}_{\text{obs}};\theta^{(l)})$. Clearly, an efficient simulation device is needed to assist the drawing of replicated datasets $\underline{\pi}^{-1(l)}_{\text{rep}}$ from the posterior predictive distribution.

Mollica and Tardella (2016, published online) also showed the usefulness of model assessment conditionally on the number m of ranked items, in order to better highlight possible deficiencies of the fitted model to reproduce specific features of the observed data. To this aim, they introduced two additional discrepancies that parallel $X^2_{(1)}$ and $X^2_{(2)}$, defined as

$$\tilde{X}^2_{(1)}(\underline{\pi}^{-1};\theta) = \sum_{m=1}^{K-1} \sum_{i=1}^{K} \frac{(r_{i,m} - r^*_{i,m}(\theta))^2}{r^*_{i,m}(\theta)} \qquad \tilde{X}^2_{(2)}(\underline{\pi}^{-1};\theta) = \sum_{m=1}^{K-1} \sum_{i<i'} \frac{(\tau_{ii',m} - \tau^*_{ii',m}(\theta))^2}{\tau^*_{ii',m}(\theta)}, \tag{9}$$

Table 3: Classification of the 23 objects included the novel **PLMIX** package.

Data manipulation	
`binary_group_ind`	`freq_to_unit`
`make_complete`	`make_partial`
`rank_ord_switch`	`unit_to_freq`
Data simulation	
`rPLMIX`	
Data description	
`paired_comparisons`	`rank_summaries`
Model estimation	
`likPLMIX`	`loglikPLMIX`
`mapPLMIX`	`mapPLMIX_multistart`
`gibbsPLMIX`	
Model selection	
`selectPLMIX`	`bicPLMIX`
Model assessment & LS adjustment	
`ppcheckPLMIX`	`ppcheckPLMIX_cond`
Datasets	
`d_apa`	`d_carconf`
`d_dublinwest`	`d_german`
`d_nascar`	

■ object of class `"function"`
■ object of class `"matrix"`

where the presence of m in the subscript refers to the evaluation of r_i and $\tau_{ii'}$ in the subset $\underline{\pi}_m^{-1} = \{\pi_s^{-1} : n_s = m\}$. The computation of $\tilde{p}_B$, obtained from (8) by adopting the generic discrepancy $\tilde{X}^2$, permits to assess the adequacy of the model estimated with the entire data set to recover the considered summary statistics within the subsets of partial orderings with the same length $m = 1, \ldots, K-1$.

Finally, similarly to the model comparison step, the LS adjustment of the posterior sample is not necessary for the posterior predictive check. This can be explained with the use of the marginal support parameters $p_i = \sum_{g=1}^{G} \omega_g p_{gi}$, which are invariant to the LS phenomenon.

5. The PLMIX package

The **PLMIX** package contains 23 objects visible to the user, classified according the their task in Table 3. There are 18 objects of class `"function"` and 5 datasets. As revealed by the overview in Table 3, the **PLMIX** package is a toolkit assisting each step of the ranking data analysis. Specifically, the available functions accomplish ranking data manipulation,

descriptive summarization, simulation, model estimation procedures, model selection criteria and model assessment diagnostics. Also the treatment of the LS issue is supported in our package. The datasets are all provided in ordering format as `"matrix"` class objects and missing positions denoted with zero entries.

The next sections illustrate in greater detail the application of the commands with real and simulated data examples.

5.1. Ranking data manipulation

Before performing a ranking data analysis, it is important to know exactly the format of the data and employ the suitable one, in order to avoid erroneous implementation or misleading interpretations of the results. As stressed in Section 5.1, the observed ordered K-tuple can be read in two different ways. Here is a toy example with $K = 4$ items, where the vector (3,1,4,2) could be regarded either as a complete *ranking* or a complete *ordering* but with a substantial different interpretation, that is,

$$\pi = (3, 1, 4, 2) \qquad \neq \qquad \pi^{-1} = (3, 1, 4, 2);$$

in fact, in the former notation π the first position is occupied by Item 2 (since $\pi(2) = 1$), whereas in the latter interpretation π^{-1} the top position is occupied by Item 3 (since $\pi^{-1}(1) = 3$). Consequently, the preliminary conversion of the data into the appropriate format is needed to start the ranking data analysis. This transformation can be performed by means of the `rank_ord_switch` function of the **PLMIX** package, switching from orderings to ranking and vice-versa for both complete and partial observations. For partial sequences (ranking/ordering), the missing positions/items have to be denoted with zero entries. For example, $\pi = (0, 1, 0, 2)$ corresponds to a top-2 ranking and can be equivalently expressed as top-2 ordering with $\pi^{-1} = (2, 4, 0, 0)$. The following instructions show the simple application of the `rank_ord_switch` routine to the first 6 partial orderings of the dataset `d_dublinwest` (Mattei and Walsh 2013), in order to convert them into the ranking format

```
> library(PLMIX)
> data(d_dublinwest)
> head(d_dublinwest)
```

```
     rank1 rank2 rank3 rank4 rank5 rank6 rank7 rank8 rank9
[1,]     7     9     4     2     8     0     0     0     0
[2,]     5     3     7     6     0     0     0     0     0
[3,]     5     7     3     0     0     0     0     0     0
[4,]     9     2     7     0     0     0     0     0     0
[5,]     3     2     0     0     0     0     0     0     0
[6,]     5     3     2     0     0     0     0     0     0
```

```
> rank_ord_switch(data=head(d_dublinwest), format="ordering")
```

```
     [,1] [,2] [,3] [,4] [,5] [,6] [,7] [,8] [,9]
[1,]    0    4    0    3    0    0    1    5    2
[2,]    0    0    2    0    1    4    3    0    0
```

```
[3,]    0    0    3    0    1    0    2    0    0
[4,]    0    2    0    0    0    0    3    0    1
[5,]    0    2    1    0    0    0    0    0    0
[6,]    0    3    2    0    1    0    0    0    0
```

The input arguments are: i) `data`, a numeric $N \times K$ data matrix of partial sequences, ii) `format`, a character string indicating the format of the input `data` argument and iii) `nranked`, an optional numeric vector of length N with the number of items ranked by each sample unit (default is `NULL`).

Another useful task is the aggregation of the replicated sequences in the observed dataset. The `unit_to_freq` routine constructs the frequency distribution of the observed sequences from the dataset of individual rankings/orderings supplied in the argument `data`. This is the output of the `unit_to_freq` when applied to the `d_german` dataset

```
> library(PLMIX)
> data(d_german)
> unit_to_freq(data=d_german)
```

```
      [,1] [,2] [,3] [,4] [,5]
 [1,]    1    2    3    4  137
 [2,]    1    2    4    3   29
 [3,]    1    3    2    4  309
 [4,]    1    3    4    2   52
 [5,]    1    4    2    3  255
 [6,]    1    4    3    2   93
 [7,]    2    1    3    4   48
 [8,]    2    1    4    3   23
 [9,]    2    3    1    4  330
[10,]    2    3    4    1   21
[11,]    2    4    1    3  294
[12,]    2    4    3    1   30
[13,]    3    1    2    4   61
[14,]    3    1    4    2   33
[15,]    3    2    1    4  117
[16,]    3    2    4    1   29
[17,]    3    4    1    2   70
[18,]    3    4    2    1   35
[19,]    4    1    2    3   55
[20,]    4    1    3    2   59
[21,]    4    2    1    3   69
[22,]    4    2    3    1   52
[23,]    4    3    1    2   34
[24,]    4    3    2    1   27
```

where the observed frequencies are indicated in the last $(K+1)$-th column. the frequency distribution helps to explore the possible presence of multiple modal patterns in the sample and

to compare the observed frequencies with those expected under specific parametric assumptions. Additionally, it can be exploited to prepare the data for the analysis with competing methods from other R packages requiring the aggregate format.

Conversely, the `freq_to_unit` function expands the frequency distribution `freq_distr` of the observed sequences into the dataset of individual rankings/orderings, as described in the following toy example with only two sequences

```
> library(PLMIX)
> K <- 5
> freq_to_unit(freq_distr=cbind(rbind(1:K,K:1), c(4,2)))

     [,1] [,2] [,3] [,4] [,5]
[1,]    1    2    3    4    5
[2,]    1    2    3    4    5
[3,]    1    2    3    4    5
[4,]    1    2    3    4    5
[5,]    5    4    3    2    1
[6,]    5    4    3    2    1
```

As preliminary ranking data manipulation, further helpful functions are the `make_partial` and the `make_complete` commands, that can be regarded specularly.

The former allows for the truncation of complete sequences according to different censoring patterns, either in a deterministic or a random way. The deterministic approach requires the user to specify the number of top positions to be retained for each sample unit in the `nranked` argument. The random approach, instead, makes use of the probabilities of top-1, top-2,..., top-$(K-1)$ censoring patterns supplied in the `probcensoring` vector to perform a stochastic truncation of the complete sequences. For example, with the following instructions one can apply a random truncation to the complete orderings of `d_german` dataset with a 60% overall rate of censored observations and equal chance of top-1 and top-2 orderings

```
> library(PLMIX)
> data(d_german)
> head(d_german)

     [,1] [,2] [,3] [,4]
[1,]    1    2    3    4
[2,]    1    2    3    4
[3,]    1    2    3    4
[4,]    1    2    3    4
[5,]    1    2    3    4
[6,]    1    2    3    4


> set.seed(57524)
> d_german_censored <- make_partial(data=d_german, format="ordering",
+   probcensoring=c(0.3, 0.3, 0.4))
```

It returns a list with two named objects given by the numeric data matrix `partialdata` of censored sequences and the numeric vector `nranked` with the number of items ranked by each sample unit after the random censoring. Here is the code to extract them and to verify the consistency of the resulting censored dataset with the probability values specified in the `probcensoring` argument

```
> head(d_german_censored$partialdata)

     [,1] [,2] [,3] [,4]
[1,]    1    2    0    0
[2,]    1    0    0    0
[3,]    1    2    0    0
[4,]    1    2    3    4
[5,]    1    2    0    0
[6,]    1    0    0    0

> N <- nrow(d_german)
> round(table(d_german_censored$nranked)/N, 2)

   1    2    4
0.30 0.29 0.41
```

The `make_partial` function is especially useful in simulation studies to investigate the impact of the censoring mechanism on the ability of the estimation procedures to recover the true generating distribution and, additionally, to verify their robustness to the censoring rate. See for example the simulation study in Mollica and Tardella (2016, published online).

Conversely, the `make_complete` function is conceived for the completion of partial orderings by filling in the missing (zero) positions with the remaining not-selected alternatives. More specifically, the completion of the partial data is performed with the random procedure determined the Plackett-Luce scheme, that is, with a sampling without replacement of the unranked items. To this aim, the positive values specified in the `probitems` argument are used as support parameters. Possible input values for the vector `probitems` could be the observed frequencies that each item has been ranked in the first position, in order to preserve the univariate feature of the observed sample. For the `d_dublinwest` dataset, this can be implemented as follows

```
> library(PLMIX)
> data(d_dublinwest)
> head(d_dublinwest)

     rank1 rank2 rank3 rank4 rank5 rank6 rank7 rank8 rank9
[1,]     7     9     4     2     8     0     0     0     0
[2,]     5     3     7     6     0     0     0     0     0
[3,]     5     7     3     0     0     0     0     0     0
[4,]     9     2     7     0     0     0     0     0     0
[5,]     3     2     0     0     0     0     0     0     0
[6,]     5     3     2     0     0     0     0     0     0
```

```
> K <- ncol(d_dublinwest)
> top_item_freq <- tabulate(d_dublinwest[,1], nbins=K)
> top_item_freq

[1]  748 3810 2300 6442 8086 2404 2370  134 3694

> set.seed(57524)
> d_dublinwest_completed <- make_complete(data=d_dublinwest, format="ordering",
+   probitems=top_item_freq)
> head(d_dublinwest_completed$completedata)

     rank1 rank2 rank3 rank4 rank5 rank6 rank7 rank8 rank9
[1,]     7     9     4     2     8     5     1     3     6
[2,]     5     3     7     6     2     9     4     1     8
[3,]     5     7     3     9     4     2     1     6     8
[4,]     9     2     7     5     4     6     3     1     8
[5,]     3     2     4     6     9     7     5     1     8
[6,]     5     3     2     4     7     6     9     1     8
```

5.2. Ranking data simulation and likelihood function

Data simulation and likelihood function are essential tasks to be suitably implemented in a model-oriented statistical package. The random generation of complete orderings is accomplished with the `rPLMIX` routine. With the following code, a random sample of $N = 5$ orderings with $K = 6$ items is drawn from a 3-component PL mixture

```
library(PLMIX)
K <- 6
G <- 3
support_par <- matrix(1:(G*K), nrow=G, ncol=K)
weights_par <- c(0.50, 0.25, 0.25)
set.seed(47201)
simulated_data <- rPLMIX(n=5, K=K, G=G, p=support_par, weights=weights_par)
```

where `support_par` is the numeric matrix with the component-specific support parameters in the rows and `weights_par` is the vector of mixture weights. In the case $G > 1$, the `rPLMIX` function returns a list of two named objects corresponding, respectively, to the vector of component memberships and to the matrix of simulated orderings, given by

```
> simulated_data$comp

[1] 1 1 2 1 3

> simulated_data$sim_data
```

```
     [,1] [,2] [,3] [,4] [,5] [,6]
[1,]    4    6    3    2    1    5
[2,]    6    5    3    4    2    1
[3,]    4    6    1    3    2    5
[4,]    5    6    4    1    2    3
[5,]    5    4    6    3    1    2
```

The calculation of the PL log-likelihood described in equation (2) is computationally intensive, especially for large data sets, since the normalization of the support parameters varies across sample units and is needed at each stage of the sequential ranking process. On the other hand, an efficient evaluation of the likelihood is crucial for the application of iterative optimization methods such as the EM algorithm, both in the MLE perspective and in the MAP estimation detailed in Section 4.2.1. Likelihood evaluation is also required for the computation of Bayesian model comparison criteria. In this regard, the `loglikPLMIX` function included in the **PLMIX** package calls a C++ routine from the R environment to reduce the computational burden. The following instructions compare the efficiency of the `loglikPLMIX` function with the `Likelihood.PL` command of the **StatRank** package for the evaluation of (2). To this purpose, first a large simulated dataset of size $N = 15000$ has been generated from the (default) uniform ranking model (corresponding to the PL with constant support parameters) and then the time needed to obtain the maximized log-likelihood value is compared

```
> library(PLMIX)
> K <- 6
> set.seed(47201)
> simulated_data <- rPLMIX(n=15000, K=K, G=1)
> PLparameters <- rep(1/K,K)
> system.time(loglikPLMIX(p=t(PLparameters), ref_order=t(1:K), weights=1,
+   pi_inv=simulated_data))

   user  system elapsed
  0.005   0.000   0.005

> library(StatRank)
> system.time(Likelihood.PL(Data=simulated_data,
+   parameter=list(m=K, Mean=PLparameters)))

   user  system elapsed
  0.181   0.002   0.182
```

Finally, the `rPLMIX` and `loglikPLMIX` functions share the `ref_order` argument relative to the *reference order* parameters of the mixture of Extended Plackett-Luce model (EPL) introduced by Mollica and Tardella (2014). The traditional PL is a special instance of the EPL with the reference order equal to the identity permutation $(1, \ldots, K)$. Since the current version of the PLMIX implements the mixture of PL models, the `ref_order` argument must be a matrix with G rows equal to the identity permutation.

5.3. Ranking data description

The **PLMIX** provides some useful utilities to conduct a preliminary exploratory analysis before fitting a ranking model.

For descriptive purposes, the `rank_summaries` command combines several sample statistics. In particular, it assists the computation of summary statistics and censoring patterns for a partial ordering/ranking dataset. The resulting list includes the following named objects:

`nranked` numeric vector with the number of items ranked by each sample unit;

`nranked_distr` the frequency distribution of the `nranked` vector;

`missing_positions` numeric vector with the number of missing positions for each item;

`mean_rank` numeric vector with the mean rank of each item;

`marginals` numeric $K \times K$ matrix of the marginal rank distributions;

`pairedcomparisons` numeric $K \times K$ paired comparison matrix.

The application of the `rank_summaries` routine to the `d_carconf` dataset returns the following information

```
> library(PLMIX)
> data(d_carconf)
> rank_summaries(data=d_carconf, format="ordering")

$nranked
  [1] 6 6 6 6 4 4 3 6 6 6 3 6 6 4 6 2 6 6 6 6 6 6 2 6 6 6 6 6 6 6 6 6 6
 [34] 3 6 6 6 6 6 6 6 6 6 3 6 6 6 6 6 4 6 4 6 3 3 4 6 6 6 3 6 4 6 6 6 6
 [67] 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 4 6 6 6 6 6 6 6 3 3 6 6
[100] 6 4 4 6 3 4 4 6 6 6 6 6 6 4 6 6 6 6 4 6 6 4 6 6 4 6 6 6 6 6 6 6 6 6
[133] 6 6 6 6 6 6 4 4 4 6 6 6 6 4 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 4 4 6 3
[166] 6 6 6 6 6 6 6 3 6 6 3 6 6 6 6 6 6 6 4 6 6 6 4 6 6 4 6 6 6 6 6 6 6
[199] 6 6 3 3 6 6 6 6 6 6 6 2 6 6 2 6 6 4 6 6 6 6 2 6 6 6 6 6 6 6 6 6 6
[232] 6 4 6 6 6 6 6 6 6 6 6 2 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 4 6 6 4 6 3 4
[265] 4 6 6 6 6 6 6 6 6 6 6 2 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 4 6
[298] 6 6 6 1 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 4 6 6 6 6 6 6 6 6 6 4 6 2 6 6 6 4
[331] 6 6 6 4 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6
[364] 6 6 6 6 6 6 6 6 6 6 4 6 6 4 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6
[397] 6 6 6 6 6 6 6 6 6 4 4 6 6 6 6 6 3 4 3 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6
[430] 6 6 4 6 4 6

$nranked_distr
Top-1 Top-2 Top-3 Top-4 Top-6
    1     8    18    43   365

$missing_positions
[1] 42 17  0 29 62 27

$mean_rank
```

```
[1] 3.559796 2.882775 3.165517 3.113300 4.493298 3.203431

$marginal_rank_distr
       Item 1 Item 2 Item 3 Item 4 Item 5 Item 6
Rank 1     86    101     87     78     28     55
Rank 2     53     87     85     86     27     96
Rank 3     46     84     78     76     46     96
Rank 4     61     74     81     69     51     72
Rank 5     57     50     62     72     74     50
Rank 6     90     22     42     25    147     39

$pairedcomparisons
       Item 1 Item 2 Item 3 Item 4 Item 5 Item 6
Item 1      0    171    179    178    250    181
Item 2    257      0    238    237    329    242
Item 3    256    197      0    230    324    226
Item 4    243    193    205      0    306    225
Item 5    148     92    111    105      0    106
Item 6    239    187    209    199    307      0
```

The command `rank_summaries` has additional logical arguments named `mean_rank`, `marginals` and `pairedcomparisons`, indicating whether the mean rank vector, the marginal rank distribution and the pairwise comparison frequencies have to be actually computed (Default is `TRUE`). The paired comparison matrix can be separately computed also with the `paired_comparisons` function.

As better detailed in Section 5.6, descriptive summaries are also involved in the model assessment step to explore compatibility between the observed dataset and specific parametric assumptions. Thus, their efficient implementation is crucial to reduce the computational time needed for the goodness-of-fit diagnostics.

5.4. Model estimation

The PL reviewed in Section 3.2 is parametrized by the item-specific quantities $\underline{p} = (p_1, \ldots, p_K)$ governing the sequential stochastic selections of the items in order of preference during the ranking process. This implies that a sample $\underline{\pi}^{-1}$ expressed in the ordering format is the natural input dataset for model inference under the PL mixture scenario. For this reason, all the functions regarding model-based ranking analysis share the `pi_inv` argument, indicating the numeric $N \times K$ matrix of observed partial orderings.

The core inferential part of the **PLMIX** package consists of the following three functions to fit a Bayesian G-component PL mixture according to the inferential procedures reviewed in Sections 4.2.1 and 4.2.2

`mapPLMIX` maximizes the posterior density via EM algorithm and returns the MAP point estimate of the PL mixture parameters;

`mapPLMIX_multistart` does the same with multiple starting values, in order to address the issue of possible local maxima in the posterior distribution;

`gibbsPLMIX` which implements the MCMC posterior simulation via Gibbs sampling, aimed at quantifying parameter uncertainty from a fully Bayesian perspective.

maximizes the posterior distribution via EM algorithm and returns the MAP point estimate of the PL mixture parameters; The above functions can be conveniently applied in a sequential way: first the MAP procedure can be launched with multiple starting values with `mapPLMIX_multistart` and then the resulting MAP estimate can be used to initialize the MCMC chain in the `gibbsPLMIX` command. Here is an example illustrating how to obtain the posterior mode for a Bayesian 3-component PL mixture fitted to the `d_apa` dataset under the noninformative prior scenario

```
> data(d_apa)
> GG <- 3
> set.seed(57524)
> outputMAP_3 <- mapPLMIX_multistart(pi_inv=d_apa, K=5, G=GG,
+   n_start=10, n_iter=400*GG, parallel=TRUE)
```

Hence, we run the EM algorithm with `n_start=10` starting values which, if not supplied by the user in the `init` argument, are randomly generated from a uniform distribution (default). The optional `centered_start` input value constraints the random starting values to be centered around the observed relative frequency that each item has been ranked first. Additionally, the `hyper` argument contains the hyperparameters values (c_{gi}, d_g and α_g) of the conjugate prior setting arranged in a list of named objects. By default, flat priors are assumed implying that the MAP estimate coincides with the MLE solution. From a computational point of view, note that the logical argument `parallel` allows to parallelize the computation and, hence, to significantly reduce the execution time.

The `mapPLMIX_multistart` automatically selects the best solution in terms of maximized posterior distribution and returns a list with the main information on the implemented MAP procedure. The MAP estimates of the component-specific support parameters and the mixture weights can be extracted from the output respectively as follows

```
> outputMAP_3$mod$P_map

           [,1]       [,2]       [,3]       [,4]       [,5]
[1,] 0.27195173 0.04679860 0.62259669 0.02287481 0.03577817
[2,] 0.08012766 0.04979918 0.02530171 0.47287656 0.37189488
[3,] 0.19123204 0.22545313 0.14637964 0.22019913 0.21673606

> outputMAP_3$mod$W_map

[1] 0.2679989 0.1316195 0.6003816
```

as well as the model-based classification of sample units into the mixture components based on the MAP posterior membership probability

```
> table(outputMAP_3$mod$classification)
```

```
   1     2     3
 466 10937  4046
```

Subsequently, we can perform an approximation of the posterior distribution by means of the GS simulation, taking advantage of the conjugate structure of the model defined in Section 4.2. The instructions are

```
> set.seed(57524)
outputGIBBS_3 <- gibbsPLMIX(pi_inv=d_apa, K=5, G=GG,
  init=list(z=binary_group_ind(outputMAP_3$mod$classification, G=GG),
  p=outputMAP_3$mod$P_map), n_iter=2100, n_burn=100)
```

In the `init` argument, the initial values are set equal to the MAP estimates. Additionally, `n_iter` and `n_burn` provide the number of total GS drawings and the length of the burn-in phase, meaning that the final posterior MCMC sample has size $L =$ `n_iter` $-$ `n_burn`. The output is a list of named objects including the posterior drawings of the supports parameters and the weights, given by

```
> outputGIBBS_3$P[1:5,]
```

```
          p1,1       p2,1      p3,1       p1,2       p2,2
[1,] 0.2749229 0.04660405 0.1921191 0.04431284 0.02666765
[2,] 0.2830427 0.04763068 0.1923358 0.04521769 0.02747666
[3,] 0.2894419 0.04997313 0.1915217 0.04554292 0.02831324
[4,] 0.2851821 0.04969482 0.1916130 0.04331265 0.02784632
[5,] 0.2819474 0.05295444 0.1861073 0.04244747 0.02865650
          p3,2      p1,3       p2,3      p3,3       p1,4
[1,] 0.2230900 0.6249881 0.01498509 0.1431151 0.02133395
[2,] 0.2191166 0.6134298 0.01588139 0.1409620 0.02228035
[3,] 0.2177711 0.6076197 0.01585510 0.1459668 0.02206100
[4,] 0.2171219 0.6157932 0.01552252 0.1462766 0.02120549
[5,] 0.2219778 0.6187978 0.01529209 0.1457384 0.02148681
          p2,4      p3,4       p1,5      p2,5      p3,5
[1,] 0.5075632 0.2224618 0.03444223 0.4041801 0.2192140
[2,] 0.5113048 0.2275917 0.03602946 0.3977065 0.2199939
[3,] 0.5047944 0.2277841 0.03533441 0.4010642 0.2169563
[4,] 0.4834107 0.2257761 0.03450659 0.4235256 0.2192124
[5,] 0.5097818 0.2276451 0.03532056 0.3933152 0.2185314
```

```
> outputGIBBS_3$W[1:5,]
```

```
            w1         w2        w3
[1,] 0.2598179 0.10005224 0.6401298
[2,] 0.2581632 0.09658437 0.6452524
[3,] 0.2646822 0.09958838 0.6357294
[4,] 0.2638245 0.10413107 0.6320445
[5,] 0.2584423 0.10540354 0.6361542
```

and the corresponding posterior likelihood and deviance function values at each iteration, obtained with

```
> outputGIBBS_3$log_lik[1:10]

 [1] -50287.50 -50289.48 -50289.90 -50291.00 -50289.84
 [6] -50290.37 -50291.42 -50291.58 -50289.18 -50287.66

> outputGIBBS_3$deviance[1:10]

 [1] 100575.0 100579.0 100579.8 100582.0 100579.7
 [6] 100580.7 100582.8 100583.2 100578.4 100575.3
```

5.5. Model comparison

The `selectPLMIX` function of the **PLMIX** package computes several Bayesian criteria to address the issue of selecting a suitable number of mixture components. Let us suppose that Bayesian PL mixtures have been fitted to the `d_apa` dataset with G varying from 1 to 3 with the code just described in Section 5.4. The comparison of the three mixtures can be performed with the following instructions

```
> outputSELECT <- selectPLMIX(pi_inv=d_apa, seq_G=1:3,
+   MAPestP=list(outputMAP_1$mod$P_map, outputMAP_2$mod$P_map,
+   outputMAP_3$mod$P_map), MAPestW=list(outputMAP_1$mod$W_map,
+   outputMAP_2$mod$W_map, outputMAP_3$mod$W_map),
+   deviance=list(outputGIBBS_1$deviance, outputGIBBS_2$deviance,
+   outputGIBBS_3$deviance))
```

The output is a list with detailed information not only on the comparison tools, but also on the separate contribution of the fitting measures and the effective number of parameters in their determination. The final values of the criteria can be obtained as

```
> outputSELECT$selection_criteria

        DIC1     DIC2    BPIC1    BPIC2    BICM1    BICM2
G=1 103204.4 103204.3 103208.3 103208.0 103233.0 103232.9
G=2 100771.9 100772.7 100779.4 100780.9 100835.6 100836.3
G=3 100591.1 100593.0 100601.3 100605.1 100685.7 100687.6
```

In this example, the decreasing trend of all the measures clearly suggests that more complex mixtures with additional components should be explored.

5.6. Model assessment

The posterior predictive check, unconditionally and conditionally on the number of ranked items, has been implemented respectively in the `ppcheckPLMIX` and `ppcheckPLMIX_cond` functions. As described in Section 4.4, the model assessment tools require the simulation of a

replicated dataset from the posterior predictive distribution for each parameter value of the GS simulation. This means that the execution time depends on both the sample sizes N and L and, hence, the computation of goodness-of-fit diagnostics is particularly time-consuming. Thanks to the combination of the R environment with the C++ code, the assessment of ranking models becomes feasible with the **PLMIX** package even for moderately large datasets. The code to perform the posterior predictive check based on the chi-square discrepancies (7) and to derive the corresponding p-values is

```
> set.seed(57524)
> outputCHECK <- ppcheckPLMIX(pi_inv=d_apa, seq_G=1:3,
+   MCMCsampleP=list(outputGIBBS_1$P, outputGIBBS_2$P,
+   outputGIBBS_3$P), MCMCsampleW=list(outputGIBBS_1$W,
+   outputGIBBS_2$W, outputGIBBS_3$W), parallel=TRUE)
> outputCHECK$post_pred_pvalue

    post_pred_pvalue_top1 post_pred_pvalue_paired
G_1                     0                 0.0000
G_2                     0                 0.6330
G_3                     0                 0.4805
```

Similarly, for the posterior predictive check based on the chi-square discrepancies (9) the script is

```
> set.seed(57524)
outputCHECKCOND <- ppcheckPLMIX_cond(pi_inv=d_apa, seq_G=1:3,
  MCMCsampleP=list(outputGIBBS_1$P, outputGIBBS_2$P,
  outputGIBBS_3$P), MCMCsampleW=list(outputGIBBS_1$W,
  outputGIBBS_2$W, outputGIBBS_3$W), parallel=TRUE)
> outputCHECKCOND$post_pred_pvalue

    post_pred_pvalue_top1_cond post_pred_pvalue_paired_cond
G_1                          0                            0
G_2                          0                            0
G_3                          0                            0
```

Remind that under correct model specification, p_B values are expected to be centered around 0.5, whereas values smaller than 0.05 are typically considered as indication of model lack-of-fit. From the results of the unconditional analysis, it follows that none of the three mixtures provides an adequate description of the observed frequencies that each item has been ranked first. Regarding the paired comparisons frequencies, instead, only the homogeneous PL is not able to recover this bivariate feature of the entire observed sample. Interestingly, the posterior predictive check conditionally on the number of ranked items reveals the inadequacy of the all estimated mixtures for both the summary statistics. This should be interpreted as an indication that a better account of the missingness mechanism is needed and, hence, a separate PL mixture analysis for each subsample of partial orderings with the same length would be preferable. See Mollica and Tardella (2016, published online) for a more in-depth analysis of the `d_apa` dataset.

5.7. Label switching adjustment

The `ppcheckPLMIX` and `ppcheckPLMIX_cond` commands can be equivalently employed to remove the possible presence of LS in the posterior MCMC sample. This step is necessary to derive meaningful point estimates of the PL mixture parameters and the related uncertainty measures. To this aim, the aforementioned functions rely on the application of the Pivotal Relabeling Algorithm (PRA) proposed by Marin *et al.* (2005), by recalling the `pra` function from the `label.switching` package (Papastamoulis 2016). The PRA acts in an intuitive way: it switches (relabels) the elements of each parameter value drawn from the posterior according to a permutation so that a certain distance from a target posterior mode is minimized. The target mode plays the role of *pivot* and can be easily identified with the MAP solution.

We simulated a sample of $N = 300$ complete orderings of $K = 4$ items from a 2-component PL mixture and the following parameter configuration

```
> K <- 4
> GG <- 2
> true_weights_param <- c(0.7, 0.3)
> true_supp_param <- rbind(c(.7,.2,.08,.02), c(.55,.3,.03,.12))
> set.seed(70476)
> simulation <- rPLMIX(n=300, K=K, G=GG, p=true_supp_param,
+   weights=true_weights_param)
> sim_orderings <- simulation$sim_data
```

The synthetic data in aggregated format are

```
> unit_to_freq(sim_orderings)
```

```
      [,1] [,2] [,3] [,4] [,5]
 [1,]    1    2    3    4   87
 [2,]    1    2    4    3   52
 [3,]    1    3    2    4   40
 [4,]    1    3    4    2    1
 [5,]    1    4    2    3   13
 [6,]    1    4    3    2    8
 [7,]    2    1    3    4   36
 [8,]    2    1    4    3   20
 [9,]    2    3    1    4    9
[10,]    2    4    1    3    5
[11,]    2    4    3    1    1
[12,]    3    1    2    4    8
[13,]    3    2    1    4    3
[14,]    3    2    4    1    1
[15,]    4    1    2    3    9
[16,]    4    1    3    2    4
[17,]    4    2    1    3    3
```

Similarly, for the posterior predictive check based on the chi-square discrepancies (9) the script is

```
> modal_orderings <- t(apply(true_supp_param,1, order, decreasing=TRUE))
> modal_orderings

     [,1] [,2] [,3] [,4]
[1,]    1    2    3    4
[2,]    1    2    4    3
```

Note that, with this support parameter setting, the component-specific modal orderings turn out to be adjacent on the ranking space in terms of the Kendall distance, since their last two positions are switched. Of course, the closeness of PL components facilitates the occurrence of LS. Subsequently, we first recorded the MAP estimate

```
> set.seed(70476)
> outputMAP <- mapPLMIX_multistart(pi_inv=sim_orderings, K=K, G=GG,
+   n_start=30, n_iter=1000, parallel=TRUE)
outputMAP$mod$P_map

          [,1]        [,2]        [,3]       [,4]
[1,] 0.6535795 0.252212857 0.061412559 0.03279511
[2,] 0.5023897 0.001954253 0.002959641 0.49269641

outputMAP$mod$W_map

[1] 0.96190009 0.03809991
```

obtained from a multiple run (`n_start` = 30) of the EM algorithm with uninformative prior densities, and then employ the posterior mode to initialize the GS. We performed a total of `n_iter` = 22000 GS iterations and discarded the first (`n_burn` = 2000) drawings as burn-in period.

```
> outputGIBBS <- gibbsPLMIX(pi_inv=sim_orderings, K=K, G=GG,
+   init=list(p=outputMAP$mod$P_map,
+   z=binary_group_ind(outputMAP$mod$classification, G=GG)),
+   n_iter=22000, n_burn=2000)
```

The traceplots of the mixture weights in Figure 1 (left) reveal the occurrence of LS. It manifests itself with step-like configurations of the traceplots indicating the transitions of the sampler from an artificial mode to another, caused by the identifiability issue. Indeed, a considerable percentage of the posterior sample is affected by LS, leading to similar (and invalid) marginal posterior means, see for example the point estimates of the support parameters given by

```
> matrix(colMeans(outputGIBBS$P), ncol=K)
          [,1]      [,2]       [,3]       [,4]

[1,] 0.5761867 0.2621617 0.07650548 0.08514619
[2,] 0.5415158 0.2795709 0.08259536 0.09631793
```

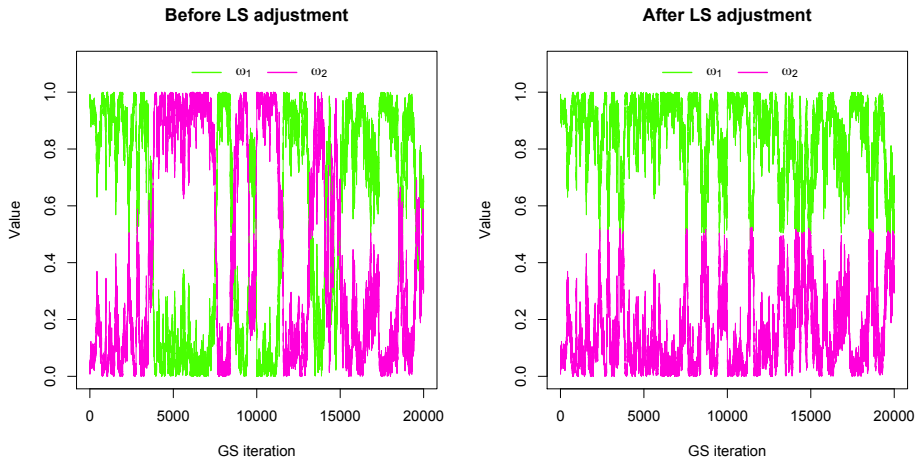

Figure 1: Traceplots of mixture weights before and after the application of the PRA to remove the LS phenomenon.

Now we can try to remove the LS phenomenon by post-processing the raw MCMC output with the PRA. To this aim, some default input values of the `ppcheckPLMIX` and `ppcheckPLMIX_cond` commands have to be modified, as indicated in the following instructions

```
> outputLS <- ppcheckPLMIX(pi_inv=sim_orderings, seq_G=GG,
+   MCMCsampleP=list(outputGIBBS$P), MCMCsampleW=list(outputGIBBS$W),
+   MAPestP=list(outputMAP$mod$P_map), MAPestW=list(outputMAP$mod$W_map),
+   label_switching_adj=TRUE, top1=FALSE, paired=FALSE, adj_post_sample=TRUE)

[1] "LABEL SWITCHING ADJUSTMENT WITH PIVOTAL RELABELING ALGORITHM"
```

The argument `label_switching_adj=TRUE` has to specified to request the adjustment of the GS simulation with the PRA, whereas `adj_post_sample=TRUE` is needed to record the relabeled posterior samples in the output list with names `final_sampleP` and `final_sampleW` for the sets of PL mixture parameters. With the logical `top1` and `paired` arguments equal to `FALSE`, the function skips the computation of the posterior predictive diagnostics. The check of the traceplots in Figure 1 (right) points out a very good performance of the PRA in removing the artificial multimodality, allowing for the computation of posterior summaries, given by

```
> matrix(apply(outputLS$final_sampleP$G_2, c(1,2), mean), ncol=K)

          [,1]      [,2]       [,3]       [,4]
[1,] 0.6742229 0.2278699 0.06266152 0.03524561
[2,] 0.4434795 0.3138627 0.09643932 0.14621852

> colMeans(outputLS$final_sampleW$G_2)

[1] 0.8511395 0.1488605
```

Finally, we can still detect a certain discrepancy between the adjusted GS estimates and the true parameter values, although the actual order of the support parameters within each group is fully recovered. As expected, when the two mixture components considerably overlap, it is more difficult to reconstruct the actual group membership of the sample units, with consequent negative effects on the final estimates. On the other hand, the performance of the GS turns out to be better than the MAP estimate corresponding to the MLE solution, since the latter completely fails to infer the minor mixture component. As testified by the MAP estimates of the mixture weights, the MLE exhibits the tendency to privilege the homogenous model, i.e., to allocate the sample units into the major mode, returning very poor estimates for the minor subpopulation.

6. Conclusions

When approaching a ranking data analysis, several issues may arise, mainly due to the peculiar structure of ranked sequences as multivariate ordinal data. First, ranking data take values in the finite discrete set of permutations, whose size $K!$ explodes with the total number of items to be ranked. In this perspective, some related difficulties are the possible occurrence of sparse data situations or the need of a manageable exploration of the ranking space when the adopted model includes permutation parameters, such as the DB and the EPL. Secondly, the presence of partial observations adds further complications. When the sample is composed of complete sequences, in fact, the PL could be estimated with methods related to the log-linear models but, in the case of partial orderings, the likelihood has to be suitably updated and the existing methods are no longer applicable. All these issues lead to computationally demanding methods and to the need of developing specialized softwares to avoid prohibitive execution time. On the other hand, this has been traditionally an obstacle for a wider use of more sophisticated models. In order to efficiently address the aforementioned practical issues and promote the effective exploitation of methodological innovations, we developed the R package **PLMIX**. This is the first package in R that implements ranking data analysis from a Bayesian inferential perspective. It relies on a hybrid code combining R and C++ and exploits the advantages of both programming languages, in particular the flexibility of the R environment and the speed of compiled C++ code to guarantee computational efficiency. It represents a comprehensive toolkit, paying special attention to each step of the model-based classification of heterogeneous partial rankings/orderings via a Bayesian PL mixture approach.

As possible directions of future work, the functions of the **PLMIX** package could be further extended for the Bayesian analysis of EPL mixtures, introduced by Mollica and Tardella (2014) in the frequentist domain, or to accommodate for the additional information provided by subject- and/or item-specific covariates. Also we note that there is little availability of R routines to handle the presence of ties in the ranking outcome and this can stimulate a further improvement. Finally, visualization techniques for the graphical illustration of ranking data analysis could be also included in a forthcoming version of the **PLMIX** package.

References

Alvo M, Yu PL (2014). *Statistical methods for ranking data.* Springer.

Ando T (2007). "Bayesian predictive information criterion for the evaluation of hierarchical Bayesian and empirical Bayes models." *Biometrika*, **94**(2), 443–458.

Babington Smith B (1950). "Discussion of Professor Ross's paper." *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **12**(1), 53–56.

Bargagliotti AE (2009). "Aggregation and decision making using ranked data." *Mathematical Social Sciences*, **58**(3), 354–366.

Biernacki C, Jacques J (2013). "A generative model for rank data based on insertion sort algorithm." *Computational Statistics & Data Analysis*, **58**, 162–176.

Bradley RA (1984). "Paired comparisons: some basic procedures and examples." In *Nonparametric methods*, volume 4 of *Handbook of Statist.*, pp. 299–326. North-Holland. doi:10.1016/S0169-7161(84)04016-5.

Bradley RA, Terry ME (1952). "Rank analysis of incomplete block designs." *Biometrika*, **39**, 324–345. ISSN 0006-3444.

Caron F, Doucet A (2012). "Efficient Bayesian inference for Generalized Bradley-Terry models." *J. Comput. Graph. Statist.*, **21**(1), 174–196. ISSN 1061-8600. doi:10.1080/10618600.2012.638220.

Clerides S, Stengos T (2012). "Love thy Neighbor, Love thy Kin: Strategy and Bias in the Eurovision Song Contest." *Ekonomia*, **15**(1), 21–44.

Cohen A, Mallows CL (1983). "Assessing goodness of fit of ranking models to data." *Journal of the Royal Statistical Society: Series D (The Statistician)*, **32**(4), 361–374. ISSN 0039-0526.

Critchlow DE, Fligner MA, Verducci JS (1991). "Probability models on rankings." *J. Math. Psych.*, **35**(3), 294–318. ISSN 0022-2496. doi:10.1016/0022-2496(91)90050-4.

Davenport A, Kalagnanam J (2004). "A computational study of the Kemeny rule for preference aggregation." In *AAAI*, volume 4, pp. 697–702.

Dempster AP, Laird NM, Rubin DB (1977). "Maximum likelihood from incomplete data via the EM algorithm." *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **39**(1), 1–38. ISSN 0035-9246. With discussion.

Diaconis PW (1987). "Spectral analysis for ranked data." *Technical Report 282*, Dept of Statistics, Stanford University.

Fligner MA, Verducci JS (1986). "Distance based ranking models." *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **48**(3), 359–369. ISSN 0035-9246.

Gelman A, Carlin JB, Stern HS, Rubin DB (2004). *Bayesian data analysis*. Second edition. Chapman & Hall/CRC, Boca Raton, FL. ISBN 1-58488-388-X.

Gelman A, Hwang J, Vehtari A (2014). "Understanding predictive information criteria for Bayesian models." *Statistics and Computing*, **24**(6), 997–1016.

Gelman A, Meng XL, Stern H (1996). "Posterior predictive assessment of model fitness via realized discrepancies." *Statistica Sinica*, **6**(4), 733–760.

Gormley IC, Murphy TB (2006). "Analysis of Irish third-level college applications data." *Journal of the Royal Statistical Society: Series A*, **169**(2), 361–379. ISSN 0964-1998. doi:10.1111/j.1467-985X.2006.00412.x.

Gormley IC, Murphy TB (2008a). "Exploring voting blocs within the Irish electorate: a mixture modeling approach." *J. Amer. Statist. Assoc.*, **103**(483), 1014–1027. ISSN 0162-1459. doi:10.1198/016214507000001049.

Gormley IC, Murphy TB (2008b). "A mixture of experts model for rank data with applications in election studies." *Ann. Appl. Stat.*, **2**(4), 1452–1477. ISSN 1932-6157. doi:10.1214/08-AOAS178.

Gormley IC, Murphy TB (2009). "A grade of membership model for rank data." *Bayesian Analysis*, **4**(2), 265–295. ISSN 1936-0975.

Gormley IC, Murphy TB (2010). "Clustering ranked preference data using sociodemographic covariates." In S Hess, A Daly (eds.), *Choice Modelling: The State-of-the-Art and the State-of-Practice: Proceedings from the Inaugural International Choice Modelling Conference*, pp. 543–569. Emerald.

Hatzinger R, Dittrich R (2012). "prefmod: An R Package for Modeling Preferences Based on Paired Comparisons, Rankings, or Ratings." *Journal of Statistical Software*, **48**(10), 1–31. ISSN 1548-7660. URL http://www.jstatsoft.org/v48/i10.

Henery RJ (1981). "Permutation probabilities as models for horse races." *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **43**(1), 86–91. ISSN 0035-9246.

Henery RJ (1983). "Permutation probabilities for Gamma random variables." *J. Appl. Probab.*, **20**(4), 822–834. ISSN 0021-9002.

Hunter DR (2004). "MM algorithms for generalized Bradley-Terry models." *Ann. Statist.*, **32**(1), 384–406. ISSN 0090-5364. doi:10.1214/aos/1079120141.

Irurozki E, Calvo Molinos B, Lozano Alonso JA (2014). "An R package for permutations, Mallows and Generalized Mallows models." *Technical report*, Dept of Computer Science and Artificial Intelligence, University of the Basque Country.

Jacques J, Biernacki C (2014). "Model-based clustering for multivariate partial ranking data." *Journal of Statistical Planning and Inference*, **149**, 201–217.

Jacques J, Grimonprez Q, Biernacki C (2014). "Rankcluster: An R package for clustering multivariate partial rankings." *The R Journal*, **6**(1), 10.

Jasra A, Holmes C, Stephens D (2005). "Markov Chain Monte Carlo methods and the label switching problem in Bayesian mixture modeling." *Statistical Science*, pp. 50–67.

Lee PH, Yu PLH (2010). "Distance-based tree models for ranking data." *Comput. Statist. Data Anal.*, **54**(6), 1672–1682. ISSN 0167-9473. doi:10.1016/j.csda.2010.01.027.

Lee PH, Yu PLH (2013). "An R package for analyzing and modeling ranking data." *BMC medical research methodology*, **13**(1), 65.

Luce RD (1959). *Individual choice behavior: A theoretical analysis*. John Wiley & Sons Inc.

Mallows CL (1957). "Non-null ranking models." *Biometrika*, **44**, 114–130. ISSN 0006-3444.

Marden JI (1995). *Analyzing and modeling rank data*, volume 64 of *Monographs on Statistics and Applied Probability*. Chapman & Hall. ISBN 0-412-99521-2.

Marin JM, Mengersen K, Robert CP (2005). "Bayesian modelling and inference on mixtures of distributions." *Handbook of Statistics*, **25**, 459–507.

Mattei N, Walsh T (2013). "PrefLib: A Library of Preference Data HTTP://PREFLIB.ORG." In *Proceedings of the 3rd International Conference on Algorithmic Decision Theory (ADT 2013)*, Lecture Notes in Artificial Intelligence. Springer.

Meilă M, Phadnis K, Patterson A (2007). "Consensus ranking under the exponential model." In R Parr, L van der Gaag (eds.), *Proceedings of the 23rd Conference on Uncertainty in Artificial Intelligence – UAI 2007*, pp. 285–294. AUAI Press. ISBN 0-9749039-3-0.

Meng XL (1994). "Posterior predictive p-values." *Ann. Statist.*, pp. 1142–1160.

Mollica C, Tardella L (2014). "Epitope profiling via mixture modeling of ranked data." *Statistics in Medicine*, **33**(21), 3738–3758. ISSN 0277-6715. doi:10.1002/sim.6224.

Mollica C, Tardella L (2016, published online). "Bayesian mixture of Plackett-Luce models for partially ranked data." *Psychometrika*. doi:10.1007/s11336-016-9530-0.

Murphy TB, Martin D (2003). "Mixtures of distance-based models for ranking data." *Comput. Statist. Data Anal.*, **41**(3-4), 645–655. ISSN 0167-9473. doi:10.1016/S0167-9473(02)00165-2.

Papastamoulis P (2016). "label. switching: An R Package for Dealing with the Label Switching Problem in MCMC Outputs." *Journal of Statistical Software*, **69**(1), 1–24.

Plackett RL (1968). "Random permutations." *Journal of the Royal Statistical Society. Series B (Methodological)*, **30**(3), 517–534.

Plackett RL (1975). "The analysis of permutations." *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, **24**(2), 193–202. ISSN 0035-9254.

Raftery AE, Satagopan Jaya M, Newton MA, Krivitsky PN (2007). "BAYESIAN STATISTICS 8." In J Bernardo, M Bayarri, J Berger, A Dawid, D Heckerman, A Smith, M West (eds.), *Proceedings of the eighth Valencia International Meeting, June 2-6, 2006*, pp. 371–416. Oxford University Press.

Soufiani HA (2014). *Revisiting Random Utility Models*. Ph.D. thesis, The School of Engineering and Applied Sciences, Harvard University.

Soufiani HA, Chen W, Parkes DC, Xia L (2013a). "Generalized method-of-moments for rank aggregation." In *Advances in Neural Information Processing Systems*, pp. 2706–2714.

Soufiani HA, Diao H, Lai Z, Parkes DC (2013b). "Generalized random utility models with multiple types." In *Advances in Neural Information Processing Systems*, pp. 73–81.

Spiegelhalter DJ, Best NG, Carlin BP, Van Der Linde A (2002). "Bayesian measures of model complexity and fit." *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **64**(4), 583–639.

Stern H (1990). "Models for distributions on permutations." *Journal of the American Statistical Association*, **85**(410), 558–564.

Stern H (1993). "Probability models on rankings and the electoral process." In *Probability models and statistical analyses for ranking data*, volume 80 of *Lecture Notes in Statist.*, pp. 173–195. Springer, New York.

Thurstone LL (1927). "A law of comparative judgement." *Psychological Reviews*, **34**, 273–286.

Vigneau E, Courcoux P, Semenou M (1999). "Analysis of ranked preference data using latent class models." *Food quality and preference*, **10**(3), 201–207.

Vitelli V, Sørensen Ø, Frigessi A, Arjas E (2014). "Probabilistic preference learning with the Mallows rank model." *arXiv preprint arXiv:1405.7945*.

Yao G, Böckenholt U (1999). "Bayesian estimation of Thurstonian ranking models based on the Gibbs sampler." *British Journal of Mathematical and Statistical Psychology*, **52**(1), 79–92.

Yellott John I J (1977). "The relationship between Luce's choice axiom, Thurstone's theory of comparative judgment, and the double exponential distribution." *J. Mathematical Psychology*, **15**(2), 109–144.

Yu PL, Lam K, Lo S (2005). "Factor analysis for ranked data with application to a job selection attitude survey." *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, **168**(3), 583–597.

Yu PLH (2000). "Bayesian analysis of order-statistics models for ranking data." *Psychometrika*, **65**(3), 281–299.

Affiliation:

Cristina Mollica
Department of Statistics
Sapienza University of Rome
Piazzale A. Moro 5
00185 Rome, Italy
E-mail: cristina.mollica@uniroma1.it

Luca Tardella
Department of Statistics
Sapienza University of Rome
Piazzale A. Moro 5
00185 Rome, Italy
E-mail: luca.tardella@uniroma1.it